

REMARKS

Application entitled to “special” status

Applicant draws attention to the extended pendency of the present application. According to MPEP 707.02, because pendency is in excess of five years, this application is entitled to “special” status.

Distinction between pull-based and push-based systems

Applicant draws attention to the fundamental distinction between two types of systems for pre-caching video assets: a “pull-based” system and a “push-based” system. In a “pull-based” system, a local server proactively fetches content and places it in the cache. In a “push-based” system, a local server passively waits for some other entity to place content in cache.

By way of analogy, in operation, a “push-based” system is somewhat like having mail delivered to your home by the Post Office. A “pull-based” system is one in which you have to go to the Post Office and ask for your mail.

The pull-based nature of the *Burns* system is apparent from the title that *Burns* selected for his patent:

PULL BASED, INTELLIGENT CACHING SYSTEM AND METHOD
FOR DELIVERING DATA OVER A NETWORK

Burns clearly describes a *pull-based* system, and *not* a push-based system.

In *Burns*, the function of causing content to be placed in the cache is carried out by each ISP **56**. In particular, this function is carried out by the combination of the pattern recognizer **116** and the scheduler **118** (see FIG. 4). These structures apparently cause content that is stored outside the ISP **56** to be copied into the ISP's local storage.

In contrast, Applicant's system provides a content manager that “pushes” content to a cache. That cache can be associated with a local server or a storage server. Applicant's system does not require that local servers or storage servers decide what video assets are to be cached.

Such decisions are carried out by a content manager that is separate from the local servers and the storage servers.

In contrast to the pull-based *Burns* system, which optimizes the inventory of video assets in a particular cache associated with a particular ISP 56, Applicant's push-based system optimizes the inventory of video assets across multiple caches, some of which are associated with local servers (such as the *Burns* ISP 56) and some of which are associated with storage servers (such as the *Burns* content server 52).

To see why *Burns* cannot optimize caches across ISPs, consider that a decision-making entity generally makes decisions on the basis of information available to it. Where the decision-making entity is the ISP 56, as is the case in *Burns*, that information is limited to locally available information, such as the content of the local storage 82, and the viewing habits of those subscribers that that particular ISP 56 serves. Under these circumstances, the best that the ISP 56 can do is to optimize the inventory of video assets in its own local storage 82. It cannot possibly optimize the inventory of video assets cached at other ISPs. For example, the ISP 56 has no way of knowing what video assets reside in caches of other ISPs. It has no way of knowing what other viewers connected to those ISPs are likely to watch, or when they are likely to watch it. A particular ISP 56 is effectively blind to much of what lies beyond it. Its tightly circumscribed world includes its local storage 78, 82, its subscribers 58, 60, and the content server 52 from which it occasionally retrieves content.

In contrast, a centralized content manager that is connected to multiple storage servers and/or local servers avoids this tunnel vision. Because the centralized content manager communicates with many local servers and storage servers, it knows exactly what is stored at each of those servers, what the subscribers to local servers are in the habit of watching, and what volume of data traffic might be expected between local servers and storage servers at any time.

Unlike the local ISP 56, Applicant's content manager sees the overall "big picture." Because of the wealth of information available to it, Applicant's content manager is not limited

to merely optimizing the disposition of video assets in a single cache associated with a single ISP that serves a limited number of subscribers. Applicant's content manager, with its broad overview of ongoing activity throughout the video delivery system, is in a far better position to optimize the disposition of video assets for the system *as a whole*.

Specific features that enable Applicant's content manager to carry out this global optimization are recited in the last paragraph of claim 73, which is parsed and annotated as follows:

a content manager

- (1) in communication with the storage servers and with the local servers,
- (2) the content manager being configured to automatically cause selected viewable data objects to be provided to servers in anticipation of expected demand for the selected viewable data objects,
- (3) the servers being selected from the group consisting of the local servers and the storage servers.

Limitation (1) recites the connectivity of the content manager to multiple servers. This connectivity provides the content manager with information from multiple servers and thereby enhances its ability to globally optimize the distribution of video assets.

Limitation (2) is intended to distinguish between a "push-based" system and a "pull-based" system. In a pull-based system, it would be the servers that would cause video assets to be provided to themselves.

Limitation (3) is intended to limit the content manager to one that delivers content to either local servers, which serve subscribers directly, and/or to storage servers, which serve either local servers or other storage servers, but not the subscribers themselves.

Section 102 rejection of claims 73, 99 and progeny

The Office appears to consider the claimed "local servers" to be met by *Burns*' ISP 56, and in particular by the continuous media server 74 and cache server 72 that are embedded in each ISP 56. The "viewer receivers" are considered to correspond to the *Burns* subscriber PC's 58, 60. The claimed "storage server" would then be the *Burns* content server 52. The claimed

“content manager,” while not expressly disclosed in *Burns*, is allegedly inherent in a system such as that disclosed by *Burns*.

Law of inherency

As best understood from MPEP 2112, in rare cases, a section 102 rejection may be made even if the cited reference fails to teach each and every limitation recited in the claim. Such a rejection is permitted only if it is clear that the missing limitation is *necessarily* present in the disclosed system. As stated by the Federal Circuit,

“To establish inherency, the extrinsic evidence must make clear that the missing descriptive matter ‘is necessary present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill. Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient’”¹

This requirement, that the missing claim element be a necessary and inevitable consequence of what was disclosed, is consistent with precedent established by the Board of Patent Appeals, which has stated that

“In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art.”²

Application of Law of Inherency to *Burns*

According to the above precedent, in order to sustain a rejection based on inherency, the disclosure of *Burns* must lead *inexorably* to the existence of a content manager that is

- (1) in communication with *both* the storage servers **52** and with the local servers **72, 74**,
- (2) configured to automatically cause selected viewable data objects to be provided to either storage servers **52** or local servers **72, 74** in anticipation of expected demand.

A necessary corollary is that if there exists even a single counter-example, i.e. if there exists the possibility that a content manager is not present in *Burns*, or that a content manager lacking even one limitation is present, a rejection based on inherency is improper.

¹ *In re Robertson*, 169 F.3d 743,743 (Fed. Cir. 1999).

² *Ex parte Levy*, 17 USPQ 2d, 1461, 1464 (Bd. Pat. App. & Inter. 1990) [emphasis in original].

The Office has cited three passages to support the proposition that a content manager *exactly as claimed* “*necessarily flows from the teaching*” of *Burns*. In particular, the Office cites

(a) col. 7, lines 30-40;

(b) col. 8, lines 23-40; and

(c) col. 5, line 64 - col. 6, line 65.

Passage (a) reads as follows:

“It is noted that both implementations of FIGS. 2 and 3 are shown and described as suitable examples for implementing various aspects of the invention. However, the network system might be implemented in a variety of arrangements. In addition, the illustrations show the subscriber units as being personal computers or work stations. However, the subscriber units can be implemented in other forms which are capable of rendering content received over the network. As examples, the subscriber computing units might include televisions, computers, game devices, handheld devices, and the like.”³

The last three sentences of passage (a) disclose different kinds of subscriber units. These are clearly irrelevant to content managers.

The first and second sentences of passage (a) do not preclude the existence of the claimed content manager. On the other hand, neither do they preclude its absence. This means that passage (a) fails to show that a content manager as claimed “*necessarily flows from the teaching*” of *Burns*.

In addition, the first and second sentences collectively state that the implementations shown in the figures are suitable for carrying out the invention. Thus, on its face, this passage already describes one embodiment that lacks the claimed content manager. This establishes that the claimed content manager is *not* a necessary constituent of the *Burns* system.

As discussed above, *Burns* discloses a pull-based system. In such a system, the ISP 56 “pulls” content from the content server 52 and places it in the cache 124. There would be no

³ *Burns*, col. 7, lines 30-40.

purpose in also having a content manager for causing content to be placed into the cache 124. After all, the ISP 56 is perfectly capable of doing this job by itself.

Having disposed of passage (a), Applicant draws attention to passage (b), which reads as follows:

The cache memory 124 serves as a quasi-temporary local storage for holding proxy copies of often used and requested target resources. The cache memory 124 can be implemented using different types of memory, including RAM storage disks (optical magnetic, etc.), and the like. If a proxy copy is stored in the cache memory 124, the target resource is served locally from the cache memory 124. If there is no proxy copy, the local service provider 110 uses the URL request to locate the target resource from a content provider and to request delivery of the target resource over the Internet. The local service provider 110 passes the target resources on to the requesting subscriber and may also cache the target resource in the cache 124 if the policy rules governing the cache are met.⁴

This passage describes a system that *reacts* to demand, rather than one that *anticipates* demand. In particular, according to the sentence:

"If there is no proxy copy, the local service provider 110 uses the URL request to locate the target resource from a content provider and to request delivery of the target resource over the Internet."

if a customer wishes to see a film, and if that film is not present in cache 124, the local service provider 110 has to go fetch the film from someplace else. This does not amount to *anticipating* demand at all. This is just *reacting* to a demand.

Moreover, the foregoing passage again describes a *pull-based* system, in which content is retrieved proactively by the local service provider 110. Passage (b) again suggests that no content manager would be needed because apparently the local service provider 110 is already quite capable of keeping the cache 124 filled with whatever it needs.

According to the Office, passage (c), which is an extended description of FIG. 2, allegedly makes clear that content is cached in anticipation of expected demand for that content. The only section of this passage that refers to the cache is the last paragraph, which essentially describes the existence of a cache for storing content requested by subscribers:

"The cache server 72 caches Internet resources, such as those requested by subscriber computers 58, 60,

⁴ Burns, col. 8, lines 23-40.

that have been downloaded from the content provider 52 to allow localized serving of those resources”⁵

This passage simply describes a cache for temporary storage of content that has apparently already been requested by subscribers. It is unclear how this suggests anything about actually *anticipating* demand, so that the content is already there before a subscriber even requests it.

The Office concedes that *Burns* never actually discloses the claimed content manager. The Office then proceeds to explain this silence by suggesting that the existence of such a content manager in effect “goes without saying.” The Office states that a content manager as recited in the claim must be an unspoken part of the system because there must be something “to pre-cache and store content frequently used.”⁶

Applicant submits a more likely explanation for *Burns*’ silence on this point. *Burns* does not describe a content manager as claimed because the *Burns* system does not need one. After all, *Burns* already has an element that places content in the cache. That element is the ISP 56 itself, and in particular the local service provider 110.

Accordingly, Applicant submits that the section 102 rejection is improper because *Burns* fails to describe or suggest the claimed content manager. Claim 99 has limitations similar to those of claim 73 and is patentable for at least the same reasons. All claims dependent on claims 73 and 99 include their respective limitations and are therefore patentable for at least the same reasons.

Section 102 rejection of claim 100 and progeny

Claim 100 recites the limitation of

“automatically managing distribution of viewable data objects among the local servers and the storage servers”

For reasons discussed in connection with claim 73, *Burns* fails to disclose automatically managing distribution *among* local servers and storage servers. At best, *Burns* discloses

⁵ *Burns*, col. 6, lines 62-65.

⁶ *Office Action*, page 3, lines 1-2.

distribution from one content server **52** to one ISP **56**. There is no indication that in *Burns*, content may be distributed from one content server **52** to another content server or from one ISP **56** to another ISP. This is particularly clear given that all *Burns* discloses is a single content server **52** and a single ISP **56**.

All that *Burns* actually discloses is that sometimes, an ISP **56** retrieves content from a content server **52** and places that content in its local storage. Such activity might amount to distribution *between* a local server and a storage server, but not to distribution *among* plural local and storage servers.

Accordingly, claim 100 and all claims dependent thereon are patentably distinct from the cited art.

Section 103 rejection of claim 121 and progeny

In rejecting claim 121, the Office merely states that “claims 121 and 122 are met by the discussions above.” Applicant notes that there are at least three pages of discussion between the discussion of claim 121 and the beginning of that portion of the office action that addresses section 103 rejections. Applicant submits that the Office’s cursory remarks are insufficient under 37 CFR 1.104(c)(2).

Nevertheless, in an effort to advance examination, Applicant speculates that the Office is somehow searching for a disclosure in which priority is attached to the video objects. Applicant urges the Office to carefully consider the language of claim 121, which states

transmitting the viewable data objects to local servers, the viewable data objects being selected in response to different priorities of the local servers for viewable data object content;

Thus the Office should be searching for disclosure of priorities attached to local servers, not to video objects. Applicant is unable to identify any such disclosure in *Burns*, which after all only discloses a single ISP **56**.

Applicant : Yvette Marie Gordon, et al.
Serial No. : 09/293,011
Filed : April 16, 1999
Page : 10 of 10

Attorney's Docket No.: 07442-009001

Accordingly, Applicant submits that *Burns* fails to teach or suggest at least the above claim limitation. Therefore, claim 121, and all claims dependent thereon, are patentably distinct over the cited art.

Summary

Now pending in this application are claims 73-127, of which claims 73, 99, 100, and 121 are independent. No additional fees are believed to be due in connection with the filing of this response. However, to the extent fees are due, or if a refund is forthcoming, please adjust our deposit account 06-1050, referencing attorney docket "07442-009001."

Respectfully submitted,

Date: May 29, 2004



Faustino A. Lichauco
Reg. No. 41,942

Fish & Richardson P.C.
225 Franklin St.
Boston, MA 02110
Telephone: (617) 542-5070
Facsimile: (617) 542-8906